Int. J. Production Economics 158 (2014) 278–289

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Management of a periodic-review inventory system using Bayesian model averaging when new marketing efforts are made

Yun Shin Lee *

KAIST College of Business, Korea Advanced Institute of Science and Technology, 85 Hoegi-Ro, Dongdaemoon-Gu, Seoul 130-722, Republic of Korea

ARTICLE INFO

Article history:
Received 28 May 2013
Accepted 30 July 2014
Available online 27 August 2014

Keywords:
Newsvendor model
Base stock policy
Critical fractile
Autocorrelated demand
Stationary demand
Non-stationary demand
Model uncertainty

ABSTRACT

Many companies invest in various marketing efforts, such as price promotion and advertising, in order to attract new customers and build customer loyalty. This paper examines the problem of setting efficient inventory levels when new marketing efforts are made and product demand is autocorrelated. We assume that the inventory manager operates with a base stock policy based on a critical fractile. If marketing has a temporary effect, the underlying demand tends to revert to a long-term equilibrium trend and the inventory manager needs to use a stationary demand model (e.g., autoregressive model) to determine the required inventory level. In contrast, if the effect is permanent, demand shocks contain an element that represents a permanent departure from previous levels and a non-stationary demand model (e.g., random walk) needs to be used instead. We show that the required inventory behaves much differently for the case of using a stationary demand model as opposed to a non-stationary model, but it is difficult in practice to identify a correct demand model in the absence of a long sampling span. In this paper, we propose an inventory model that explicitly acknowledges uncertainty over stationary and non-stationary demand models in response to new marketing efforts. The proposed model averages the inventory policies of the two demand models, weighted by each model's posterior probability. This is an extension of Bayesian model averaging. Simulation results demonstrate that the Bayesian model averaging inventory model improves the inventory system.

1. Introduction

Many companies invest in various marketing efforts, such as price promotion and advertising, in order to attract new customers and build customer loyalty. Since the early 1970s, price promotion has emerged as an important part of the marketing mix (Srinivasan et al., 2004). In June 2011, promotions represented 39% grocery sales in the UK (Retailgazette, 2011). Furthermore, during January–June 2011, the total measured advertising expenditure in the US reached $71.5 billion. The advertising spending in companies like AT&T Inc. and Procter & Gamble Co. amounts to billions of dollars per year (KantarMedia, 2011). There is no doubt that new marketing efforts are frequently made with significant investment. It is therefore an important and practical problem to determine the required inventory stock of a consumer product for which new marketing efforts are made.

New marketing actions (single or multiple) induce a series of unexpected movements (shocks) in demand as a result of six marketing factors: instantaneous effects, delayed response, purchase reinforcement, performance feedback, decision rules, and competitive reactions (Dekimpe and Hanssens, 1995a). Dekimpe and Hanssens (1995a, 1999) and Hanssens (1998) conclude that the total over-time effectiveness of these movements on underlying consumer demand will be either temporary or permanent. On one hand, some of these movements are temporary, in that after a number of periods the underlying demand process reverts to a long-term equilibrium trend, such as a fixed mean or an upward/downward trend. In this case, the underlying demand process is defined as stationary. Other changes are permanent if a portion of these changes is carried forward and sets a new trend in performance. There is a permanent departure from pre-expenditure performance levels and the underlying demand process becomes non-stationary.

When setting inventory levels in response to new marketing efforts, it is critical to correctly interpret the nature of marketing effects on product demand, as it determines the level of expected demand and the magnitude of demand uncertainty during the replenishment lead time. If the nature of demand shocks is temporary, the future reversion of demand has a direct effect on anticipated demand levels and also reduces demand uncertainty during the lead time. The stronger the reversion, the lower the uncertainty in the lead time demand (LTD), which may, in turn, necessitate lower safety stock levels. The inventory manager then needs to choose a stationary demand model to set inventory

* Tel.: +82 2958 3339; fax: +82 2958 3359.
E-mail address: yunshin@business.kaist.ac.kr

http://dx.doi.org/10.1016/j.ijpe.2014.08.016
0925-5273/

levels. In contrast, when marketing has a persistent effect, the underlying demand process diverges over time by indefinitely accumulating demand shocks and necessitates a higher level of safety stock to mitigate demand uncertainty. In this case, a non-stationary demand model needs to be used instead.

To illustrate, a sales promotion may persuade a thousand consumers to switch to a product at the promotional price during each sale period (Dekimpe and Hanssens, 1995a). If these consumers return to their previous purchasing habits once the promotion has ended, the resulting fluctuations are temporary in nature and have no impact upon the underlying consumer demand trend. The inventory manager must use a stationary demand model and be careful not to overinterpret short-run demand fluctuations as an indication of future demand patterns and set correspondingly high inventory quantities. In contrast, if two hundred of the promotion-captured customers not only make an initial purchase but also continue to purchase the product in future, the demand shocks have a persistent effect and we would see sales deviating permanently from pre-promotional levels. In this scenario, the inventory manager is required to set a higher safety stock level by using a non-stationary demand model in order to buffer against persistent demand shocks.

In this paper, we assume that the inventory manager operates with a base stock (order-up-to) policy based on a critical fractile (e.g., Graves, 1999) and no backorders are assumed. Under this policy, one orders a variable quantity every fixed period of time so that an inventory position is maintained at a predefined base stock level. We further assume that the underlying demand process is autocorrelated (Urban, 2005; Charnes et al., 1995; Chen and Blue, 2010). A frequent practice is then to make inventory decisions on the assumption that the true demand model in response to new marketing efforts is known with certainty. An overconfident inventory manager, believing her knowledge of the nature of demand fluctuations to be accurate, chooses either a stationary or non-stationary demand model to estimate inventory base stock levels. However, in the absence of a long sampling span, it is difficult in practice to capture the long-term trend and distinguish the correct nature of demand shocks. Often, small samples are falsely thought to represent the properties of the statistical process that generated them. This is known as the “law of small numbers” (Camerer, 1989; Rabin, 2002).

Rather than employing one or other of the demand model assumptions by default, one may take a step further and use a statistical test, such as a unit root test, as a formal criterion for making the distinction between stationary and non-stationary demand processes (e.g., Nijs et al., 2001; Pauwels et al., 2002). When sample size is small, it is again difficult to choose a correct demand model, as conventional unit root tests have low statistical power in a finite sample (Diebold and Rudebusch, 1991; DeJong et al., 1992). To sum up, the underlying demand model cannot be identified with certainty using both contextual expertise and a statistical rationale in small samples. Furthermore, it may well be costly to wait for more periods to pass and obtain more data in order to identify the trend more clearly. However, as the required inventory levels behave much differently for one demand model compared to the other, the incorrect demand model results in the under or overestimation of inventory levels, leading to increased inventory costs.

We propose an inventory policy that directly incorporates the inherent uncertainty over stationary and non-stationary demand models in response to new marketing efforts, by using Bayesian model averaging. Bayesian model averaging is a complete Bayesian solution to average over possible models. The concept of Bayesian model averaging was introduced by Leamer (1978), and has recently received significant attention in the statistics and econometrics literature, in particular from Raftery et al. (1997), Hoeting et al. (1999), and Raftery and Zheng (2003). We assume that one of the stationary or non-stationary models is the true demand model once new marketing efforts are introduced, but that we do not know which it is. Starting from a prior about which model is true and observing demand, we compute the posterior probabilities that each is the true model by applying Bayes' theorem. We then average over the inventory decisions made by the two models, weighted by each model's posterior probability. In this paper, structural results of the proposed inventory model are also discussed. Specifically, the Bayesian model averaging inventory model estimates consistent order quantities based on a critical fractile and provides better performance, as measured by a logarithmic scoring rule, than using any single model.

The paper in hand relates to several studies that use a Bayesian framework to deal with parameter uncertainty for specific demand models (see, e.g., Azoury, 1985; Azoury and Miyaoka, 2009; Lovejoy, 1990). Using Bayes' theorem, the unknown parameter is periodically updated based on newly obtained demand observations. Yet these inventory models cannot address uncertainty about the structure of the underlying demand generating model, i.e., demand model uncertainty. In this paper, we use the Bayesian framework to update the belief about candidate demand models on the basis of past observations to explicitly account for model uncertainty, which in our case arises from uncertainty about the nature of demand fluctuations after new marketing efforts. While non-parametric approaches (e.g., Bookbinder and Lordahl, 1989; Levi et al., 2007) are established to negate the need to make assumptions about the demand model, they are limited to independent demand processes and cannot be applied to serially correlated demand processes. A semi-parametric approach in Lee (2014) provides consistent estimates of the critical fractile independently of a forecasting model if the demand process follows a stationary autoregressive demand process and the forecasting model is within the autoregressive integrated class. Unlike the non-parametric and semi-parametric inventory models, our proposed Bayesian model averaging (BMA) inventory model enables us to deal with model uncertainty in independent, serially correlated, as well as non-stationary demand processes. As such, our paper can be seen as a first step towards suitably modifying and adapting the recent developments in the Bayesian model averaging method seen in the statistics and econometrics literature to the practical problem facing inventory management, namely that of setting inventory levels in response to new marketing efforts.

The interaction between two functional areas, marketing and operations, is recurrently discussed in the literature. See, for example, Tang (2010), Ma et al. (2013), and Marques et al. (2014). We contribute to this type of literature focusing on the issue of marketing efforts and ordering decisions. In particular, our work is linked to the literature that addresses the classical single-period inventory problem with advertising, where advertising stimulates the demand. Khouja and Robbins (2003) assume that the mean demand is both increasing and concave in advertising expenditure (i.e., the returns of advertising have a diminishing effect on sales) and demand variance is also a function of advertising expenditure. They obtain the optimal advertising expenditure and ordering quantity that maximizes the expected profit or the probability of achieving a target profit. Their model assumes that the demand process is independent and the effect of advertising on the underlying demand is known. Lee and Hsu (2011) and Guler (2014) recently extend this model to the distribution-free newsboy problem. In contrast, our model considers autocorrelated demand and the effect of marketing actions on the mean and variance of demand is characterized by the autocorrelation parameter. In most practical situations, we shall indeed observe autocorrelation in the demand process, especially when new marketing efforts are made